arXiv:1603.07813v1 [cs.SI] 25 Mar 2016

Chatty Maps: Constructing sound maps of urban areas from social media data

Luca Maria Aiello[1], Rossano Schifanella[2], Daniele Quercia[3], and Francesco Aletta[4]

[1]Yahoo Labs, London, UK
[2]University of Turin, Italy
[3]Bell Labs, Cambridge, UK
[4]University of Sheffield, UK

Abstract

Urban sound has a huge influence over how we perceive places. Yet, city planning is concerned mainly with noise, simply because annoying sounds come to the attention of city officials in the form of complaints, while general urban sounds do not come to the attention as they cannot be easily captured at city scale. To capture both unpleasant and pleasant sounds, we applied a new methodology that relies on tagging information of geo-referenced pictures to the cities of London and Barcelona. To begin with, we compiled the first urban sound dictionary and compared it to the one produced by collating insights from the literature: ours was experimentally more valid (if correlated with official noise pollution levels) and offered a wider geographic coverage. From picture tags, we then studied the relationship between soundscapes and emotions. We learned that streets with music sounds were associated with strong emotions of joy or sadness, while those with human sounds were associated with joy or surprise. Finally, we studied the relationship between soundscapes and people's perceptions and, in so doing, we were able to map which areas are chaotic, monotonous, calm, and exciting.Those insights promise to inform the creation of restorative experiences in our increasingly urbanized world.

1 Introduction

Studies have found that long-term exposure to urban noise (in particular, to traffic noise) results into sleeplessness and stress [19], increased incidence of learning impairments among children [51], and increased risk of cardiovascular morbidity such as hypertension [53] and heart attacks [21, 48].

Since those health hazards are likely to reduce life expectancy, a variety of technologies for noise monitoring and mitigation have been developed over the years. However, those solutions are costly and do not scale at the level of an entire city. City

officials typically measure noise by placing sensors at a few selected points. They do so mainly because they have to comply with the Environmental Noise Directive (END) [52], which requires the management of noise levels only from specific sources, such as road traffic, railways, major airports and industry. To fix the lack of scalability of a typical solution based on sensors, in distinct fields, researchers have worked on ways of making noise pollution estimation cheap. They have worked, for example, on epidemiological models to estimate noise levels from a few samples [31], on capturing samples from smartphones or other pervasive devices [28, 47, 29, 8, 32], and on mining geo-located data readily available from social media (e.g., foursquare, twitter) [22].

All this work has focused, however, on the negative side of urban sounds. Pleasant sounds have been left out from the urban planning literature, yet they have been shown to positively impact city dwellers' health [34, 4]. Only a few researchers have been interested in the whole "urban *soundscape*". In the World Soundscape Project[1], for example, composer Raymond Murray Schafer and colleagues defined soundscape for the first time as *"an environment of sound (or sonic environment) with emphasis on the way it is perceived and understood by the individual, or by a society"* [44]. That early work eventually led to a new International Standard, ISO 12913, where soundscape is defined as *"[the] acoustic environment as perceived or experienced and/or understood by a person or people, in context"* [23]. Since that work, there remains a number of unsolved challenges though.

First, there is no shared vocabulary of urban sounds. Back in the early days of the World Soundscape Project, scholars collected sound-related terms and provided a classification of sounds [44], but that classification was meant to be neither comprehensive nor systematic. Signal processing techniques for automatically classifying sounds have recently used labeled examples [43, 42], but, again, those training labels are not organized in any formal taxonomy.

Second, studying the relationship between urban sounds and people's perceptions is hard. So far the assumption has been that a good proxy for perceptions is noise level. But perceptions depend on a variety of factors; for example, on what one is doing (e.g., whether (s)he is at a concert). Therefore, policies focusing only on the reduction of noise levels might well fall short.

Finally, urban sounds cannot be captured at scale and, consequently, they are not considered when planning cities [1]. That is because the collection of data for managing urban acoustic environments has mainly been relegated to small-scale surveys [20] [45, 46, 15].

To partly address those challenges, we used geo-referenced social media data to map the soundscape of an entire city, and related that mapping to people's emotional responses. We did so by extending previous work that captured urban smellscapes from social media [38] with four main contributions:

- We collected sound-related terms from different online and offline sources and arranged those terms in a taxonomy. The taxonomy was determined by matching the sound-related terms with the tags on 1.8 million geo-referenced Flickr pictures in Barcelona and London, and by then analyzing how those terms co-occured across the pictures to obtain a term classification (co-occuring terms are

[1] `www.sfu.ca/~truax/wsp.html`

expected to be semantically related). In so doing, we compiled the first urban sound dictionary and made it publicly available: in it, terms are best classified into 6 top-level categories (i.e., transport, mechanical, human, music, nature, indoor), and those categories closely resemble the manual classification previously derived by aural researchers over decades.

- Upon our picture tags, we produced detailed sound maps of Barcelona and London at the level of street segment. By looking at different segment types, we validate that, as one expects, pedestrian streets host people, music, and indoor sounds, while primary roads are about transport and mechanical sounds.

- For the first time, we studied the relationship between urban sounds and emotions. By matching our picture tags with the terms of a widely-used word-emotion lexicon, we determined people's emotional responses across the city, and how those responses related to urban sound: fear and anger were found on streets with mechanical sounds, while joy was found on streets with human and music sounds.

- Finally, we studied the relationship between a street's sounds and the perceptions people are likely to have of that street. Perceptions came from *soundwalks* conducted in two cities in UK and Italy: locals were asked to identify sound sources and report them along with their subjective perceptions. Then, from social media data, we determined a location's expected perception based on the sound tags at the location.

2 Methodology

The main idea behind our method was to search for sound-related words (mainly words reflecting potential sources of sound) on geo-referenced social media content. To that end, we needed to get hold of two elements: the sound-related words, and the content against which to match those words.

2.1 Sound Words

We obtained sound-related words from the most comprehensive research project in the field – the World Soundscape Project – and from the most popular crowd-sourced online repository of sounds – Freesound.

Schafer's Words

The World Soundscape Project is an international research project that initiated the modern study of acoustic ecology. In his book "The Soundscape", the founder of the project, R. Murray Schafer, coined the term soundscape and emphasized the importance of identifying pleasant sounds and using them to create healthier environments. He described how to classify sounds, appreciating their beauty or ugliness, and offered exercises (e.g., "soundwalks") to help people become more sensitive to sounds. An

entire chapter was dedicated to the classification of urban sounds, which was based on literary, anthropological, and historical documents. Sounds were classified into six categories: natural, human, societal (e.g., domestic sounds), mechanical, quiet, and indicators (e.g., horns and whistles). Our work used that classification as well: to associate words with each category, three annotators independently hand-coded the book's sections dedicated to the category, and the intersection of the three annotation sets (which is more conservative than the union) was considered, resulting in a list of 236 English terms.

Crowdsourced Words

Freesound is the largest public online collaborative repository of audio samples: 130K sounds annotated with 1.5M tags are publicly available through an API. Out of the unique tags (which were 65K), we considered only those that occurred more than 100 times (the remaining ones were too sparse to be useful), resulting in 2.2K tags, which still amounted to 76% of the total volume as the tag frequency distribution was skewed. However, those tags covered many topics (including usernames, navigational markers, sound quality descriptions, and synthesized sound effects) and reflected ambiguous words at times (e.g., "fan" might be a person or a mechanical device) and, as such, needed to be further filtered to retain only words related to sounds or physical sound sources. One annotator manually performed that filtering, which resulted into a final set of 229 English terms.

In addition to that set of words, there is an online repository specifically tailored to urban sounds called Favouritesounds[2]. This site hosts crowdsourced maps of sounds for several cities in the world: individuals upload recordings of their favorite sounds, place them on the city map, and annotate them with free-text descriptions. By manually parsing the 6K unique words contained in those descriptions, we extracted 243 English terms.

2.2 Geo-referenced content

Having two sets of sound-related words at hand, we needed social media data against which those words had to be matched. 17M Flickr photos taken between 2005 and 2015 along with their tags were made publicly available in London and Barcelona. In those two cities, we identified each *street segment* from OpenStreetMap[3] (OSM is a global group of volunteers who maintain free crowdsourced online maps). We then collated tags in each segment together by considering the augmented area of the segment's polyline, an area with an extra space of 22.5 meters on each side to account for positioning errors typically present in geo-referenced pictures [38, 37].

We found that, among the three crowdsourced repositories, Freesound words matched most of the picture tags and offered the widest geographic coverage (Figure 1), in that, they matched 2.12M tags and covered 141K street segments in London. Also,

[2] `favouritesounds.org`

[3] A segment is often a street's portion between two road intersections but, more generally, it includes any outdoor place (e.g., highways, squares, steps, footpaths, cycle-ways).

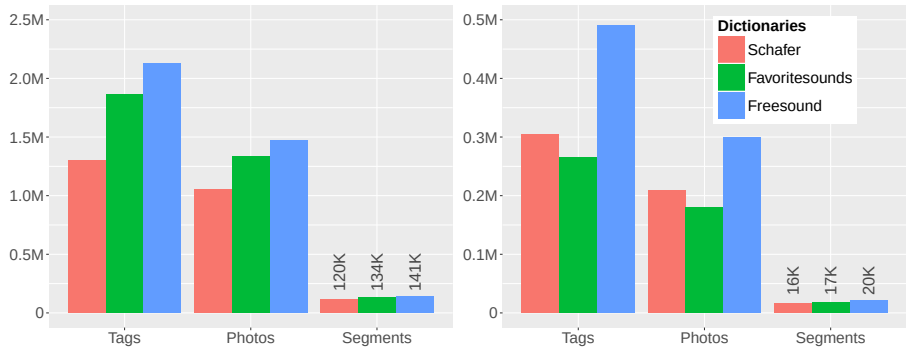

Figure 1: Coverage of the three urban sound dictionaries. Number of tags, photos, street segments that had at least one smell word from each vocabulary in Barcelona and London. Each bar is a smell vocabulary. Shafer was extracted from Shafer's book "The Soundscape", while the other two were online repositories. The best coverage was offered by FreeSound.

Freesound's words offered a far better coverage than what Shafer's did, with a broad distribution of tags over street segments (Figure 2).

Since the words of the other online repository considerably overlapped with Freesound's (67% Favoritesounds tags are also in Freesound), we worked only with Freesound (to ensure effective coverage) and with Shafer's classification (to allow for comparability with the literature).

2.3 Categorization

To discover similarities, contrasts, and patterns, sound words needed to be classified. Schafer already did so. Our Schafer's words are classified into 7 main categories - *nature*, *human*, *society*, *transport*, *mechanical*, *indicators*, and *quiet* - and each category might have a subcategory (e.g., society includes the subcategories *indoor* and *entertainment*). By contrast, Freesound's words are not classified. However, by looking at which Freesound's words co-occur in the same locations, we could discover similarities (e.g., nature words could co-occur in parks, while transport words in trafficked streets). The use of community detection to extract word categories had been successfully tested in previous work that extracted categories of smell words [38]. Compared to other clustering techniques (e.g., LDA [10], K-means [27]), a community detection technique has the advantage of being fully non-parametric and quite resilient to data sparsity. Therefore, we applied it here as well. We first built a co-occurrence network where nodes were Freesound's words, and undirected edges were weighted with the number of times the two words co-occurred in the same Flickr pictures as tags. The semantic relatedness among words naturally emerged from the network's community structure: semantically related nodes ended up being both highly clustered together and weakly connected to the rest of the network. To determine the clustering structure, we could have used any of the literally thousands of different community detection

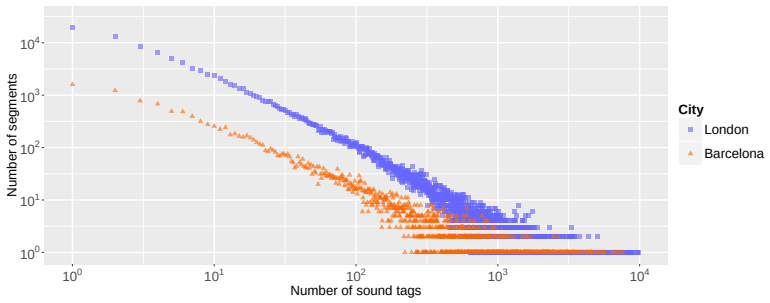

Figure 2: Number of street segments (y-axis) containing a given number of picture tags that match Freesound terms (x-axis) in London and Barcelona. Many streets had a few tags, and only a few streets have a massive number of them. London has $141K$ segments with at least one tag (and 15 tags in each segment, on average), Barcelona $20K$ (25 tags per segment on average).

Figure 3: Urban sound taxonomy. Top-level categories are in the inner circle; second-level categories are in the outer ring; and examples of words are in the outermost ring. For space limitation, in the wheel, only the first categories (those in the inner circle) are complete, while subcategories and words represent just a sample.

algorithms that have been developed in the last decade [17]. None of them always returns the "best" clustering. However, since Infomap had shown very good performance across several benchmarks [17], we opted for it to obtain the initial partition of our network [41]. Infomap's partitioning resulted in many clusters containing semantically-related words, but it also resulted in some clusters that were simply too big to possibly be semantically homogeneous. To further split those clusters, we applied the community detection algorithm by Blondel *et al.* [12], which has been found to be the second best performing algorithm [17]. This algorithm stops when no "node switch" between communities increases the overall *modularity* [33], which measures the overall quality of the resulting partitions[4]. The result of those two steps is the grouping of sound words in hierarchical categories. Since a few partitions of words could have been too fine-grained, we manually double-checked whether this was the case and, if so, we merged all those sub-communities that were under the same hierarchical partition and that contained strongly-related sound words.

Figure 3 sketches the resulting classification in the form of a sound wheel. This wheel has six main categories (inner circle), each of which has a hierarchical structure with variable depth from 0 to 3. For brevity, the wheel reports only the first level fully (inner circle), while it reports samples for the two other levels. Despite spontaneously emerging from word co-occurrences and being fully data-driven, the classification in the wheel strikingly resembles Schafer's. The three categories *human*, *nature*, and *transport* are all in both categorizations. The category *quiet* is missing because it does not match any tag in Freesound, as one would expect. The remaining categories are all present but arranged at a different level: *music* and *indoor* are at the first level in the wheel, while they are at the second level in Schafer's categorization; the *mechanical* category in the wheel collates two of Schafer's categories into one: *mechanical* and *indicator*.

Freesound not only offered a classification similar to Schafer's and to recent working groups' classifications [13, 43] (speaking to its external validity) but also offered a richer vocabulary of words. By looking at the fraction $\frac{tag_c}{tag}$ of sound words in category c that matched at least one geo-reference picture tag (tag_c) over the total number of tags in the city (tag), we saw that Freesound resulted in a full representation of all sound categories (Figure 4), while Shafer's resulted in a patchy representation of many categories. Therefore, given its effectiveness, Freesound was chosen as the sound vocabulary for the creation of the urban sound wheel. Only the wheel's top-level categories were used. The full taxonomy is, however, available online[5] for those wishing to explore specialized aspects of urban sounds (e.g., transport, nature).

3 Validation

With our sound categorization, we were able to determine, for each street segment j, its *sound profile* $sound_j$ in the form of a 6-element vector. Given sound category c, the

[4] If one were to apply Blondel's right from the start, the resulting clusters would be less coherent than those produced by our approach.

[5] `http://goodcitylife.org/chattymaps/project.html`

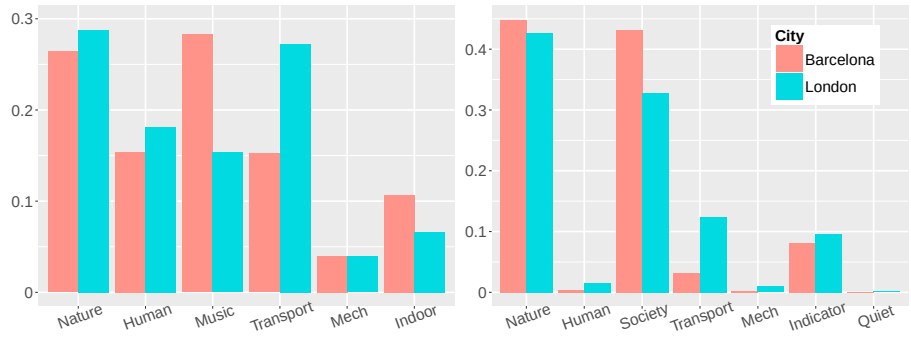

Figure 4: Fraction $\frac{tag_c}{tag}$ of picture tags that matched sound category c over all the tags in the city.

element $sound_{j,c}$ is:

$$sound_{j,c} = \frac{tag_{j,c}}{tag_j}; \tag{1}$$

where $tag_{j,c}$ is the number of tags at segment j that matched sound category c, and tag_j is the total number of tags at segment j. To make sure the sound categories c's we had chosen resulted into reasonable outcomes, we verified whether different street types were associated with sound profiles one would expect (Section 3.1), and whether those profiles matched official noise pollution data (Section 3.2).

3.1 Street types

One way of testing whether the 6-category classification makes sense in the city context is to see which pairs of categories do not tend to co-occur spatially (e.g., nature and transport should be on separate streets). Therefore, for each street segment, we computed the pairwise Spearman rank correlation ρ between the fraction of sound tags in category c_1 and that of sound tags in category c_2, across all segments (Figure 5). That is, we computed $\rho_j(sound_{j,c_1}, sound_{j,c_2})$ across all j's. We found that the correlations were either zero or negative. This meant that the categories were either orthogonal (i.e., the categories of *human*, *indoor*, *music*, *mechanical* show correlations close to zero) or geographically sorted in expected ways (with a $\rho = -0.50$, *nature* and *transport* are seen, on average, on distinct segments).

To visualize the geographic sorting of sounds, we marked each street segment with the sound category that had the highest z-score in that segment (Figure 6). The z-scores reflect the extent to which the fraction of sound tags in category c at street segment j deviated from the average fraction of sound tags in c at all the other segments:

$$z_{sound_{j,c}} = \frac{sound_{j,c} - \mu(sound_c)}{\sigma(sound_c)}; \tag{2}$$

where $\mu(sound_c)$ and $\sigma(sound_c)$ are the mean and standard deviation of the fractions of tags in sound category c across all segments. We then reported the most prominent

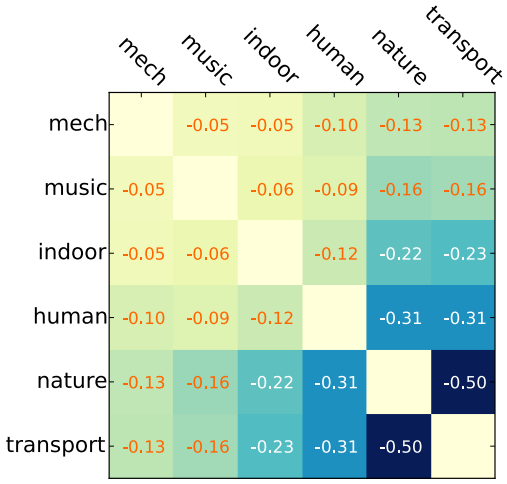

Figure 5: Pairwise rank correlations between the fraction of sound tags in category c_1 ($sound_{j,c_1}$) and the faction in category c_2 ($sound_{j,c_2}$) across all segments j's in London.

sound at each street segment in Figure 6: *traffic* was associated with street junctions and main roads, *nature* with parks or greenery spots, and *human* and *music* with central parts or with pedestrian streets.

One indeed expects that different street types (Table 1 reports the most frequent types in OSM) would be associated with different sounds. To verify that, we computed the average z-score of a sound category c for the segments with street type t:

$$\bar{z}_{sound_c,type_t} = \frac{\sum_{j \in S_t}(z_{sound_{j,c}})}{|S_t|}; \tag{3}$$

where S_t is the set of segments of (street) type t. Figure 7 reports the average values of those z-scores. Each clock-like representation refers to a street type, and the sound categories unfold along the clock: positive (negative) z-score values are marked in green (red) and suggest a presence of a sound category higher (lower) than the average one. By looking at the positive values , we saw that primary, secondary and tertiary streets (which contain cars) were associated with transport sounds; construction sites with mechanical sounds; footways and tracks (often embedded in parks) were associated with nature sounds; residential and pedestrian streets were associated with human, music, and indoor sounds. Then, by looking at the negative values, we learned that primary, secondary, tertiary, and construction streets were not associated with nature; and the other street types were not associated with sounds related to transport.

3.2 Noise pollution

The most studied aspect of urban sounds is the issue of noise pollution. Despite the importance of that issue, there is no reliable and high-coverage noise measurement data for world-class cities. There is a great number of participatory sensing applications

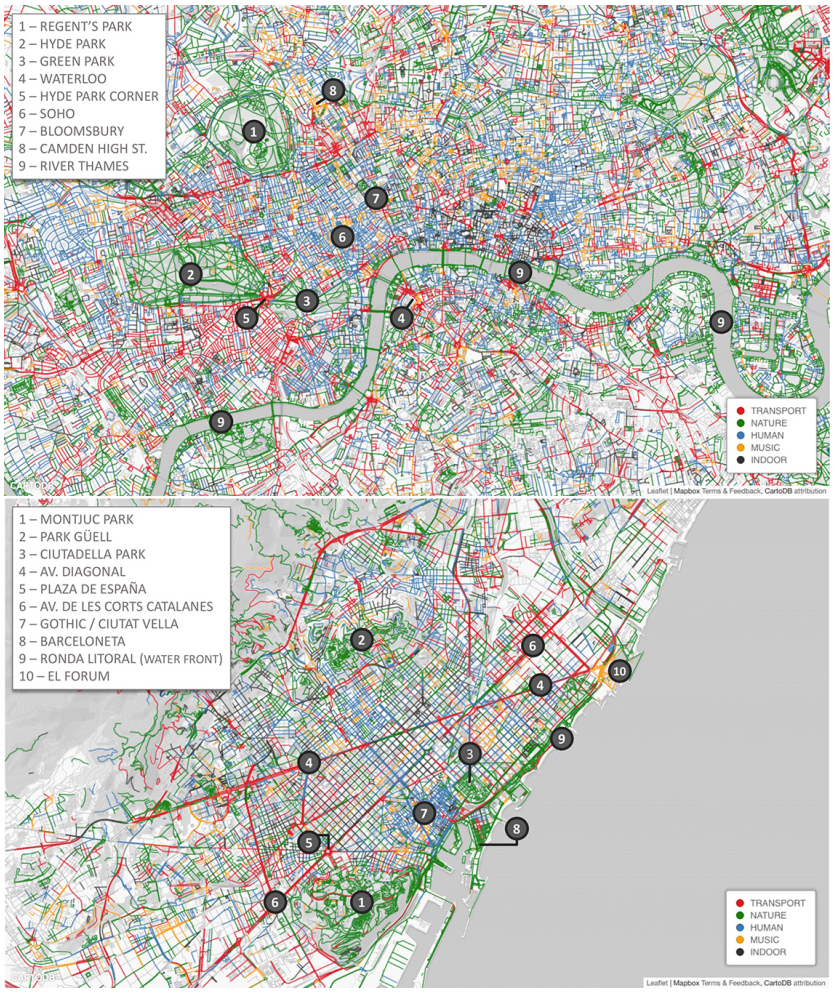

Figure 6: Urban Sound Maps of London (top) and Barcelona (bottom). Each street segment is marked with the sound category c that has the highest z-score for that segment ($z_{sound_{j,c}}$). In London, natural sounds are found in Regent's Park(1); Hyde Park(2); Green Park(3); and all around the River Thames(9). By contrast, transport sounds are around Waterloo station(4) and on the perimeter of Hyde Park (5). Human sounds are found in Soho(6) and Bloomsbury(7), and music is associated with the small clubs on Camden High Street(8). In Barcelona, natural sounds are found in Montjuic Park(1); Park Guell(2); and Ciutadella Park(3), and on the beaches of Barceloneta(8) and Ronda Litoral(9). By contrast, annoying and chaotic sounds are found on the main road of Avinguda Diagonal(4), on Plaza de Espana(5), and on Avinguda De Les Corts Catalanes(6). Human sounds are found in the historical center called Gothic/Ciutat Vella(7), and music in the open-air arena of El Forum(10). Only segments with at least 5 sound tags were considered.

Street type	Description
Footway	Designated footpaths mainly or exclusively for pedestrians. This includes walking tracks and gravel paths.
Residential	Roads that serve as an access to housing, without function of connecting settlements. Often lined with housing.
Pedestrian	Roads used mainly or exclusively for pedestrians in shopping and residential areas. They may allow access of motorised vehicles only for very limited periods of the day.
Track	Roads for mostly agricultural or forestry uses. Tracks are often rough with unpaved surfaces.
Primary	A major highway linking large towns, normally with 2 lanes not separated by a central barrier.
Secondary	A highway which is not part of a major route, but nevertheless forming a link in the national route network, normally with 2 lanes.
Tertiary	Roads connecting smaller settlements or roads connecting minor streets to more major roads.
Construction	Active road construction sites. Major road and rail construction schemes that typically require several years to complete.

Table 1: Description of the eight most frequent street types in Open Street Map.

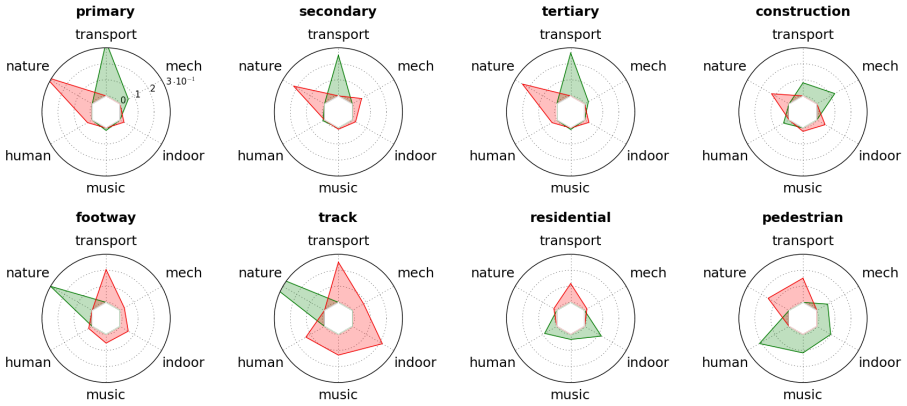

Figure 7: Average z-scores of the presence of six sound categories for segments of each street type ($\overline{z}_{sound_c,type_t}$). Positive values are in green, and negative ones are in red. The first clock-like representation refers to primary roads and shows that transport sounds have z-score of 0.3 (they are present more than average), while nature sounds have z-score of -0.3 (they are present less than average). The number of segment per type ranges between 1K and 25K, with the only exception of the "construction" type that has only 83 segments. Confidence intervals around the average values range between 10^{-2} and 10^{-3}.

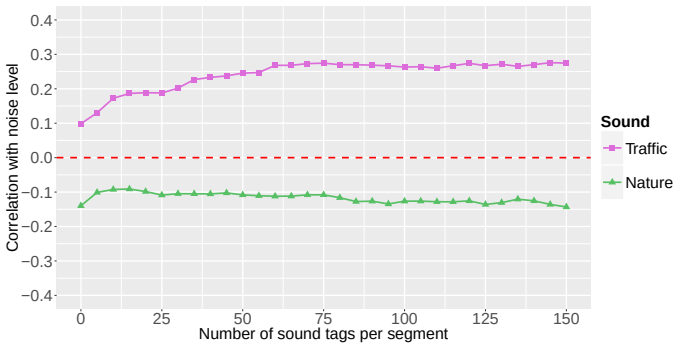

Figure 8: Spearman correlation between the fraction of tags at segment j that matched category c ($sound_{j,c}$) and j's noise levels (expressed as Equivalent Weighted Level values in dB) as the number of tags per street segment (x-axis) increases. All correlations are statistically significant at the level of $p < 0.01$.

that manage databases of noise levels in several cities, and some of them are publicly accessible [28, 47, 29, 8, 32], but all of them offer a limited geographic coverage of a city.

Barcelona is an exception, however. In 2009, the city council started a project, called Strategic Noise Map, whose main goal was to monitor noise levels and ultimately find new ways of limiting sound pollution. The project has a public API[6] that returns noise values at the level of street segment for the whole city. For each segment, we collected the four dB values provided: three yearly averages for the three times of the day (day, from 7am to 9pm; evening, from 9pm to 11pm; and night, from 11pm to 7am), and one aggregate value, the Equivalent Weighted Level (EWL), that averages those three values adding a 5dB penalty to the evening period, and a 10dB to the night period. With a practice akin to the one used for air quality indicators [16, 9], those noise level values are estimated by a prediction model that is bootstrapped with field measurements [18]. In the case of Barcelona, the model is bootstrapped with 2.300 short-span noise measurements lasting at most 15 minutes, usually taken during daytime, and with 100 long-span ones lasting from 24 hours to a few days.

To see whether noise pollution was associated with specific sound categories, we considered the street segments with at least N tags and computed, across all the segments, the Spearman rank correlations $\rho_j(EWL_j, sound_{j,c})$ between segment j's EWL values (in dB) and j's fraction of picture tags that matched category c[7] (Figure 8). The idea was to determine not only which categories were associated with noise pollution but also how many tags were needed to have a significant association. We found that noise pollution was positively correlated ($p < 0.01$) with *traffic* ($0.1 < \rho < 0.3$) and negatively correlated with *nature* ($-0.1 < \rho < -0.2$), and those results did hold for low

[6]Interactive Map of Noise Pollution in Barcelona `http://w20.bcn.cat:1100/WebMapaAcustic/mapa_soroll.aspx?lang=en`

[7]In computing the correlations, we used the method by Clifford *et al.* [14] to addresses spatial autocorrelations.

values of N, suggesting that only a few hundred tags were needed to build a representative sound profile of a street.

4 Emotional and Perceptual Layers

Sounds can be classified in ways that reflect aspects other than semantics - they may be classified according to, for example, their emotional qualities or the way they are perceived. Therefore, we now show how social media helps extracting the emotional layer (Section 4.1) and the perceptual layer (Section 4.2) of urban sound.

4.1 Emotional Layer

Looking at a location through the lens of social media makes it possible to characterize places from different points of views. Sound has a highly celebrated link with emotions, especially music sound [25, 56], and it has a considerable effect on our feelings and our behavior.

One way of extracting emotions from geo-referenced content is to use a word-emotion lexicon known as EmoLex [30]. This lexicon classifies words into eight primary emotions: it contains binary associations of 6,468 terms with the their typical emotional responses. The 8 primary emotions (anger, fear, anticipation, trust, surprise, sadness, joy, and disgust) come from Plutchik's psychoevolutionary theory [36], which is commonly used to characterize general emotional responses. We opted for EmoLex instead of other commonly used sentiment dictionaries (such as LIWC [35]) as it made it possible to study finer-grained emotions.

We matched our Flickr tags with the words in EmoLex and, for each street segment, we computed its *emotion profile*. The profile consisted of all Plutchik's primary emotions, in that, each of its elements was associated with an emotion:

$$emotion_{j,e} = \frac{tag_{j,e}}{tag_j}; \tag{4}$$

where $tag_{j,e}$ is the number of tags at segment j that matched primary emotion e. We then computed the corresponding z-score:

$$z_{emotion_{j,e}} = \frac{emotion_{j,e} - \mu(emotion_e)}{\sigma(emotion_e)}. \tag{5}$$

By computing the Spearman rank correlation $\rho_j(z_{sound_{j,c}}, z_{emotion_{j,e}})$, we determined which sound was associated with which emotion. From Figure 9, we see that joyful words were associated with streets typically characterized by music and human sounds, while they were absent in streets with traffic. Traffic was, instead, associated with words of fear, anticipation, and anger. Interestingly, words of sadness (together with those of joy) were associated with streets with music, words of trust with indoors, and words of surprise with streets typically characterized by human sounds.

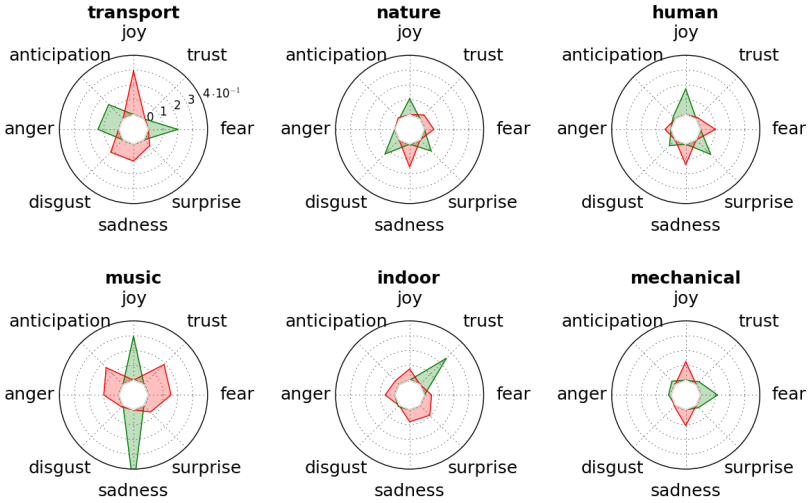

Figure 9: Correlation between $z_{sound_{j,c}}$ and $z_{emotion_{j,e}}$. Each clock-like representation refers to a sound category. The different emotions unfold around the clock, and the emotions that are associated with the sound category are marked in green (positive correlations) or in red (negative emotions). All correlations are statistically significant at the level of $p < 0.01$.

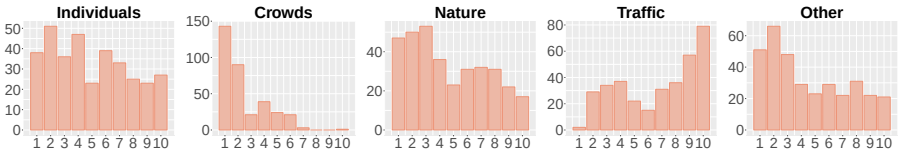

Figure 10: Frequency distributions of the survey's scores for sound presence (from 1 to 10) across categories: individuals, crowds, nature, traffic, and other. Sounds of individuals are scored in the full 1-to-10 range, while sounds of crowds are typically scored with a value of 1 or 2 as they might have been absent most of the times.

4.2 Perceptual Layer

From our social media data, we knew the extent to which a potential source of sound was present on a street. If we knew how people usually perceived that source as well, we could have estimated how the street was likely to be perceived.

One way of determining how people usually perceive sounds in the city context is to run *soundwalks*. These were introduced in the late 60s [50] and are still common among acoustic researchers nowadays [49, 24]. Therefore, to determine people's perceptions, one of the authors conducted *soundwalks* across 8 areas in Brighton & Hove (UK) and 11 areas in Sorrento (Italy) in April and October. They involved 37 participants (UK: 16 males, 5 females, $\mu_{age} = 38.6$, $\delta_{age} = 11.5$; Italy: 10 males, 6 females, $\mu_{age} = 34.7$, $\delta_{age} = 7.1$) with a variety of backgrounds (e.g., acousticians, architects, plan-

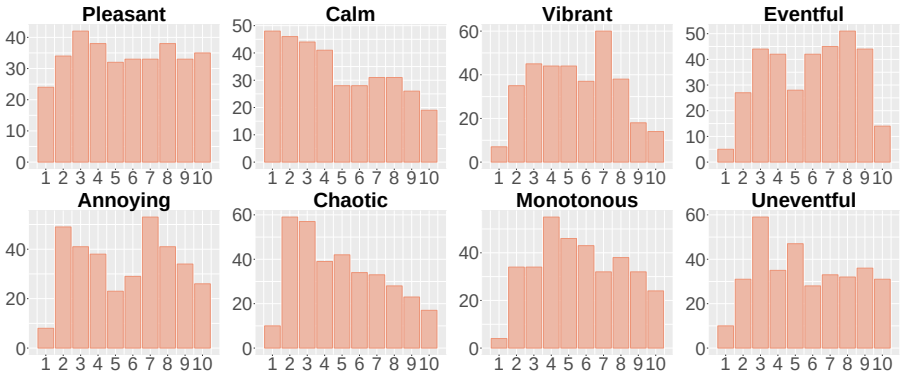

Figure 11: Frequency distributions of the survey's perception scores (from 1 to 10) for each perception category. Most of the perceptions are scored in the full 1-to-10 range.

ning professionals, local authorities, and environmental officers). The experimenter led the participants along a pre-defined route and stopped at selected locations. At each of the locations, participants were asked to listen to the acoustic environment for two minutes and to complete a structured questionnaire (Table 2) inquiring about sound sources' noticeability [6], soundscape attributes [6], overall soundscape quality [6, 26] and soundscape appropriateness [5]. The questionnaire classified urban sounds into five categories (traffic, individuals, crowds, nature, other) as it is typically done in soundwalks [3, 2], and the perceptions of such sounds into eight categories (pleasant, chaotic, vibrant, uneventful, calm, annoying, eventful, and monotonous, after Axelsson *et al.* [7]'s work).

Those soundwalks resulted into 342 tuples, each of which represents a participant's report about sounds and perceptions at a given location. Each tuple had thirteen [1,10] values: five values reflecting the extent to which the five sound categories were reported to be present, and the other eight reflecting the extent to which the eight perceptions were reported. More technically, $sound_{k,c}$ is the score for sound category c at tuple k, and $perception_{k,f}$ is the score for perception category f at tuple k. The frequency distributions of $sound_{k,c}$ (Figure 10) suggest that the participants experienced both streets with only a few sounds, and streets with many. Also, they rarely experienced crowds and came across traffic and, only at times, nature. Instead, the frequency distributions of $perception_{k,f}$ (Figure 11) suggest that the participants experienced streets with very diverse perceptual profiles, resulting in the use of the full [1,10] score range for all perceptions.

To see which sounds participants tended to experience together, we computed the rank cross-correlation $\rho_k(sound_{k,c_1}, sound_{k,c_2})$ (left panel of Figure 13). Amid crowds, the participants reported high score in the category 'individuals'. These two sound categories –individuals and crowds– had similar sound profiles so much so that the category 'crowds' could be experimentally replaced by the category 'individuals' in the specific instance of those soundwalks. Furthermore, as one would expect, the presence of traffic was associated with the absence of individuals, crowds, and nature.

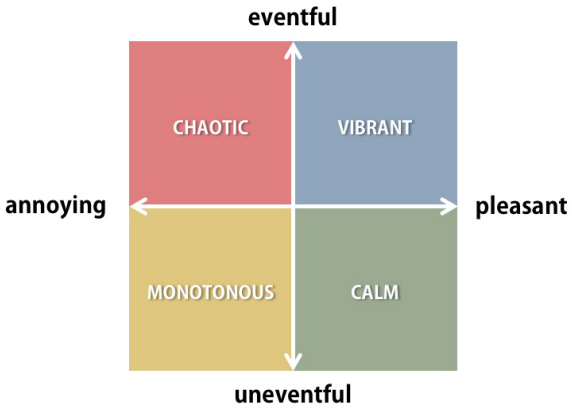

Figure 12: Two principal components describing how study participants perceived urban sound. The combination of the first component "uneventful *vs.* eventful" with the second component "annoying *vs.* pleasant" results in four main ways of perceiving urban sounds: vibrant, calm, monotonous, and chaotic.

To then see which perceptions participants tended to experience together, we computed the rank cross-correlation $\rho(perception_{k,f_1}, perception_{k,f_2})$ (right panel of Figure 13). Perceptions meant to have opposite meanings indeed resulted into negative correlations (pleasant *vs.* annoying, eventful *vs.* uneventful, vibrant *vs.* monotonous, and calm *vs.* chaos). Interestingly, with their near-zero correlation, pleasantness and eventfulness were orthogonal - when a place was eventful, nothing could have been said about its pleasantness.

To see which sounds participants experienced together with which perception, we computed the rank correlation $\rho_k(sound_{k,c}, perception_{k,f})$ (left panel of Figure 14). On average, vibrant areas tended to be associated with crowds, pleasant areas with individuals, calm areas with nature, and annoying and chaotic areas with traffic. In a similar way, Axelsson *et al.* studied the principal components of their perceptual data [7] and found very similar results: they found that two components best explain most of the variability in the data (Figure 12).

Finally, to map how streets are likely to be perceived, we needed to estimate a street's *expected* perception given the street's sound profile. The sound profiles came from our social media data, while the expected perception could have been computed from our soundwalks' data. We had already computed the correlations between sounds and perceptions (left panel of Figure 14). However, those correlations are not expected values (accounting for, e.g., whether a perception is frequent or rare) but they simply are strength measures. Therefore, we computed the probability of perception f given sound category c as:

$$p(f|c) = \frac{p(c|f) \cdot p(f)}{p(c)} \tag{6}$$

To compute the composing probabilities, we needed to discretize our [1,10] values

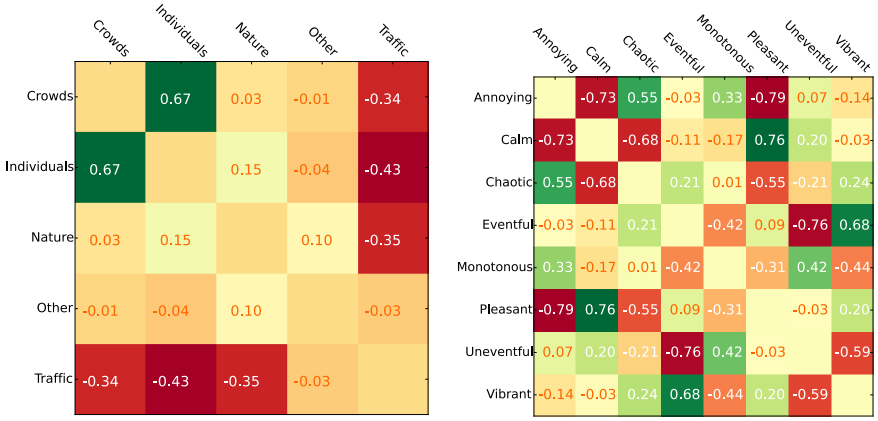

Figure 13: Pairwise rank cross-correlations between the survey's sound scores $sound_{k,c}$ (left) and its perception scores $perception_{k,e}$ (right).

taken during the soundwalks, and did so by segmenting them into quartiles. We then computed:

$$p(c|f) = \frac{Q_4(c \wedge f)}{Q_4(f)} \tag{7}$$

$$p(c) = \frac{Q_4(c)}{Q_4(c^*)} \quad ; \quad p(f) = \frac{Q_4(f)}{Q_4(f^*)} \tag{8}$$

where $Q_4(c)$ is the number of times the sound category c occurred in the fourth quartile of its score; $Q_4(c^*)$ is the number of times *any* sound occurred in its fourth quartile; and $Q_4(c \wedge f)$ is the number of times sound c as well as perception f occurred in their fourth quartiles.

The conclusions drawn from the resulting conditional probabilities (right panel in Figure 14) did not differ from those drawn from the previously shown sound-perception correlations (left panel). As opposed to the correlation values, none of the conditional probabilities were very high (all below 0.33). This is because the conditional probabilities were estimated through the gathering of perceptual data in the wild[8] and, as such, the mapping between perception and sound did not result in fully-fledged probability values. Those values are best interpreted not as raw values but as ranked values. For example, nature sounds were associated with calm only with a probability 0.34, yet calm is the strongest perception related to nature as it ranks first.

The advantage of conditional probabilities over correlations is that they offer principled numbers that are properly normalized and could be readily used in future studies. They could be used, for example, to draw an aesthetics map, a map that reflects the emotional qualities of sounds. In the maps of Figure 15, we associated each segment with the color corresponding to the perception with the highest value

[8]It has been shown that, in soundwalks, perception ratings are affected by not only sounds but also visual cues (e.g., greenery has been found to modulate soundscape 'tranquillity' ratings [55, 54]).

Question	Items	Scale extremes (1-10)
To what extent do you presently hear the following five types of sounds?	Traffic noise (e.g., cars, trains, planes), sounds of individuals (e.g., conversation, laughter, children at play), crowds of people (e.g., passers, sports event, festival), natural sounds (e.g., singing birds, flowing water, wind in the vegetation), other noise (e.g. sirens, construction, industry)	[Do not hear at all, ..., Dominates completely]
Overall, how would you describe the present surrounding sound environment?	—	[Very bad, ..., Very good]
Overall, to what extent is the present surrounding sound environment appropriate to the present place?	—	[Not at all, ..., Perfectly]
For each of the 8 scales below, to what extent do you agree or disagree that the present surrounding sound environment is...	pleasant, chaotic, vibrant, uneventful, calm, annoying, eventful, monotonous	[Strongly disagree, ..., Strongly agree]

Table 2: The questionnaire used during the soundwalk. For each question, participants could express their preference on a 10-point ordinal scale.

of $p_j(f) = \sum_c p(f|c) \cdot p_j(c)$, where $p_j(c) = sound_{j,c}$, which is the fraction of tags at segment j that matched sound category c. $p_j(f)$ is effectively the probability that perception f is associated with street segment j, and the strongest f is associated with j. By mapping the probabilities of sound perceptions in London (top panel of Figure 15) and Barcelona (bottom panel), we observed that trafficked roads were chaotic, while walkable parts of the city were exciting. More interestingly, in the soundscape literature, monotonous areas have not necessarily been considered pleasant (they fall into the annoying quadrant of Figure 12), yet the beaches of Barcelona were monotonous (and rightly so), but might have been pleasant as well.

5 Discussion

A project called SmellyMaps mapped urban smellscapes from social media [38], and this work - called ChattyMaps - has three main similarities with it. First, the taxonomy of sound and that of smell were both created using community detection algorithms, and both closely resembled categorizations widely used by researchers and practitioners in the corresponding fields. Second, the ways that social media data was mapped onto streets (e.g., buffering of segments, use of long/lat coordinates on the pictures) is the same. Third, in both works, the validation was done with official data (i.e., with air quality data and noise pollution data). However, the two works differ as well, and they do so in three main ways. First, as opposed to SmellyMaps, ChattyMaps studied a variety of urban layers: not only the urban sound layer but also the emotional, perceptual, and sound diversity layers. Second, smell words were derived from smellwalks (as no other source was available), while sound words were derived from the online platform

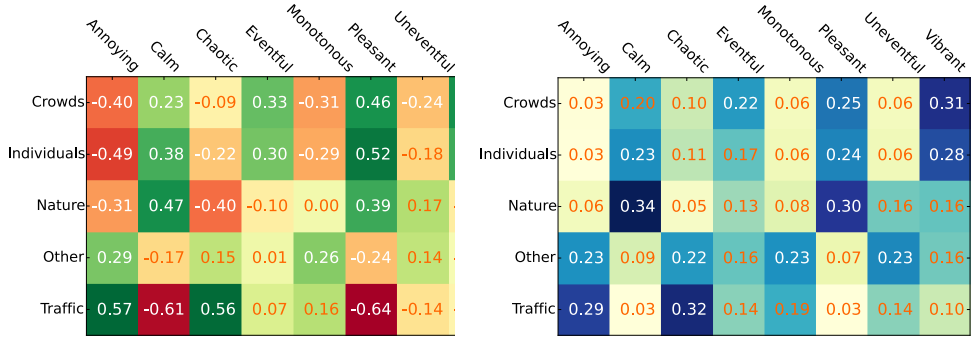

(a) Correlations between the survey's sound scores $sound_{k,c}$ and its perception scores $perception_{k,e}$. Sounds of crowds, for example, are perceived to be pleasant and vibrant but not annoying.

(b) Probability $p(f|c)$ that perception f was reported at a location with sound category c.

Figure 14: Relationship between sounds and perceptions in the soundwalk survey data.

of Freesound. Third, since SmellyMaps showed that picture tags were more effective than tweets in capturing geographic-salient information, ChattyMaps entirely relied on Flickr tags.

Our approach comes with a few limitations, mainly because of data biases. The urban soundscape is multifaceted: the sounds we hear and the way we perceive them change considerably with small variations of, for example, space (e.g., simply turning a corner) and time (e.g., day *vs.* night). By contrast, social media data has limited resolution and coverage, and that results into false positives. At times, sound tags do not reflect real sounds because of either misannotations or the figurative use of tags (picture Figure 16(*a*)). Fortunately, those cases occur rarely. By manually inspecting 100 photos with sound tags, no false positive was found: 87 pictures were correctly tagged, and 13 referred to sounds that were plausible yet hard to ascertain.

Even when tags refer to sounds likely present in an area, they might do so partially. For example, the tags on the picture of Figure 16(*b*) consisted of traffic terms (rightly) but not of nature terms, and that was a partial view of that street's soundscape. This risk shrinks as the number of sound tags for the segment increases. Indeed, let us stick with the same example: Figure 16(*c*) was taken few meters away from (*b*), and its tags consisted of nature terms.

To partly mitigate noise at boundary regions, we did two things. First, as described in Section 2, we added a buffer of 22.5 meters around each segment's bounding box. This has been commonly done in previous work dealing with geo-referenced digital content [37, 38]. It is hard to measure automatically how many tags are needed to get high confidence sound profiles, but we estimated it to be around 20-25 tags (Figure 8), if official air quality data is used for validation.

Second, we associated sound *distributions* (and not individual sounds) with street segments. The 6-dimensional sound vector was normalized in [0,1] to have a proba-

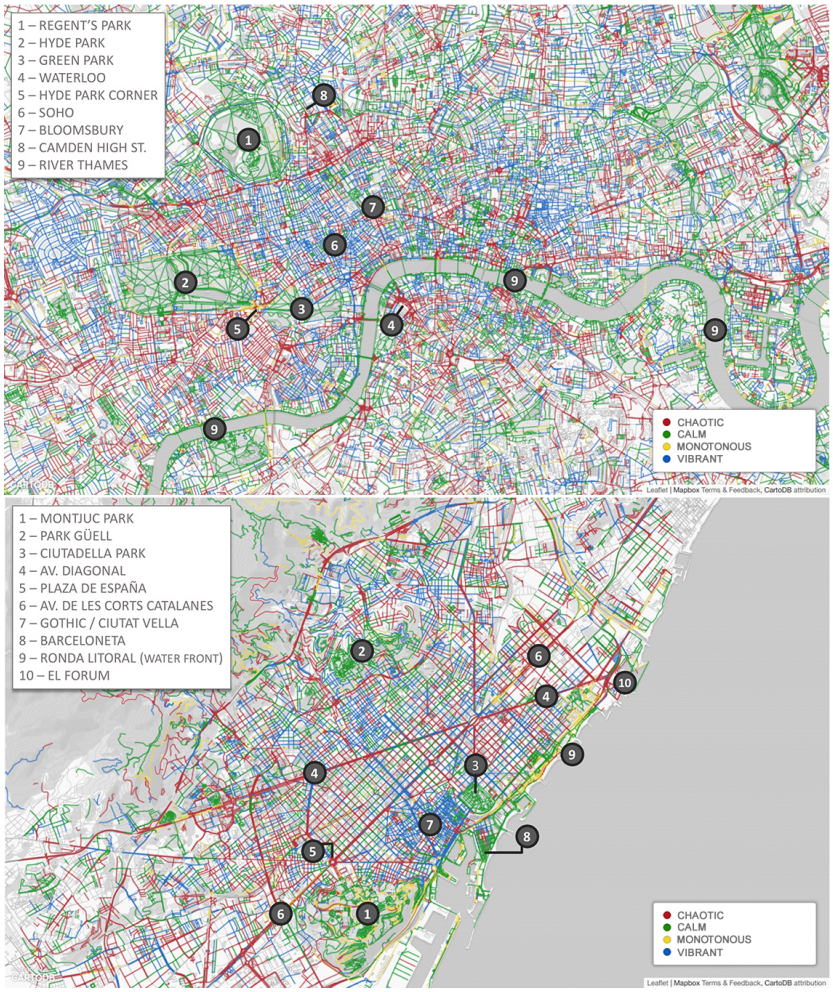

Figure 15: Perceptual maps of London (top) and Barcelona (bottom). At each segment, the perception f with the highest probability was reported (i.e., with the highest $p_j(f)$). In London, calm sounds were found in Regent's Park(1); Hyde Park(2); Green Park(3); and all around the River Thames(9). By contrast, chaotic sounds were around Waterloo station(4) and Hyde Park corner (5). Vibrant sounds were found in Soho(6), Bloomsbury(7), and Camden High Street(8). In Barcelona, calm sounds were found in Montjuic Park(1); Park Guell(2); and Ciutadella Park(3), and on the beach of Barceloneta(8). By contrast, on the beach in front of Ronda Litoral(9), we found monotonous sounds. Chaotic sounds were found on the main road of Avinguda Diagonal(4), on Plaza de Espana(5), and on Avinguda De Les Corts Catalanes(6). Vibrant sounds were found in the historical center called Gothic/Ciutat Vella(7), and a bit in the open-air arena of El Forum(10), which was also characterized by chaotic sounds.

Figure 16: Examples of ambiguously tagged pictures. *(a)* street art in Brick Lane tagged with the term "screaming", and the same location Carriage Drive with Hyde Park tagged with opposing terms related to *(b)* traffic sounds and *(c)* nature sounds.

bilistic interpretation. In Figures 16(b) and (c), nature sounds were predominant, yet traffic-related sounds varied from 20% to 2% depending on the different parts of that street.

More generally, to have a more comprehensive view of this phenomenon, we determined each segment's sound diversity by computing the Shannon index:

$$diversity_j = -\sum_c sound_{j,c} \cdot ln(sound_{j,c}) \tag{9}$$

where $sound_{j,c}$ is the fraction of tags at segment j that matched sound category c. After removing zero diversity values (often associated with segments having only one tag, which made 28% of segments in Barcelona, and 35% segments in London), we saw that the frequency distribution of diversity (Figure 17, left) had two peaks in 1 (for both cities) and in 1.5 for London and in 2.0 for Barcelona. Then, by mapping those values (Figure 18), we saw that the values close to the first peak were associated with parks and suburbs, and those close to the second peak (and higher) were associated with the central parts of the two cities. Furthermore, the diversity did not depend on the number of tags per segment and became stable for segments with at least 10 tags (Figure 17, right).

6 Conclusion

We showed that social media data makes it possible to effectively and cheaply track urban sounds at scale. Such a tracking was effective because the resulting sounds were geographically sorted across street types in expected ways, and they matched noise pollution levels. The tracking was also cheap because it did not require the creation of any additional service or infrastructure. Finally, it worked at the scale of an entire city, and that is important, not least because, before our work, there had been *"nothing in sonography corresponding to the instantaneous impression which photography can create ... The microphone samples details and gives the close-up but nothing corresponding to aerial photography."* [44]

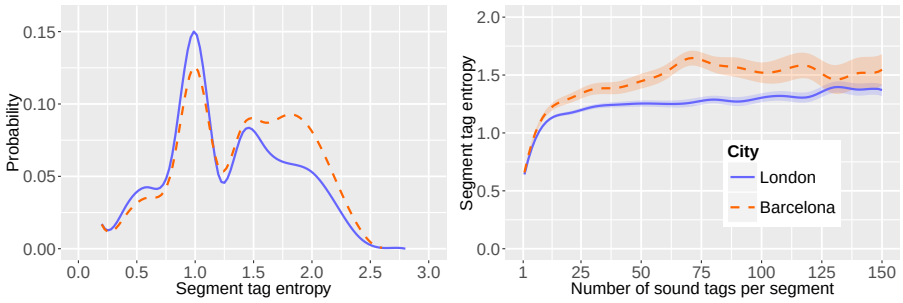

Figure 17: Diversity (entropy) of sound tags. Frequency distribution (left), and how the diversity varies with the number of tags per street segment (right). Segments with zero diversity (28% in Barcelona, 35% in London) were excluded.

However, whereas landscapes can be static, soundscapes are dynamic [11]. Their perceptions are affected by demography (e.g., personal sensitivity to noise, age), context (e.g., city layout), and time (e.g., day *vs.* night, weekdays *vs.* weekends). Future studies could partly address those issues by collecting additional data and by comparing models of urban sounds generated from social media with those generated from GIS techniques.

Nonetheless, no matter what data one has, fully capturing soundscapes might well be impossible. Our work has focused on identifying potential sonic events. To use a food metaphor, those events are the raw ingredients, then the aural architecture (which comes with the acoustic properties of trees, buildings, streets) is the cooking style, and the soundscape is the dish [11].

To unite hitherto isolated studies in a new synergy, in the future, we will conduct a comprehensive multi-sensory research of cities, one in which visual [40, 39], olfactory [38], and sound perceptions are explored together.

The ultimate goal of this work is to empower city managers and researchers to *"find solutions for an ecologically balanced soundscape where the relationship between the human community and its sonic environment is in harmony"*, as Schafer famously (and prophetically) remarked in the late 70s [44].

Acknowledgments

We thank the Barcelona City Council for making the noise pollution data available. F. Aletta received funding through the People Programme (Marie Curie Actions) of the European Union's 7^{th} Framework Programme FP7/2007-2013 under REA grant agreement *n* 290110, SONORUS "Urban Sound Planner".

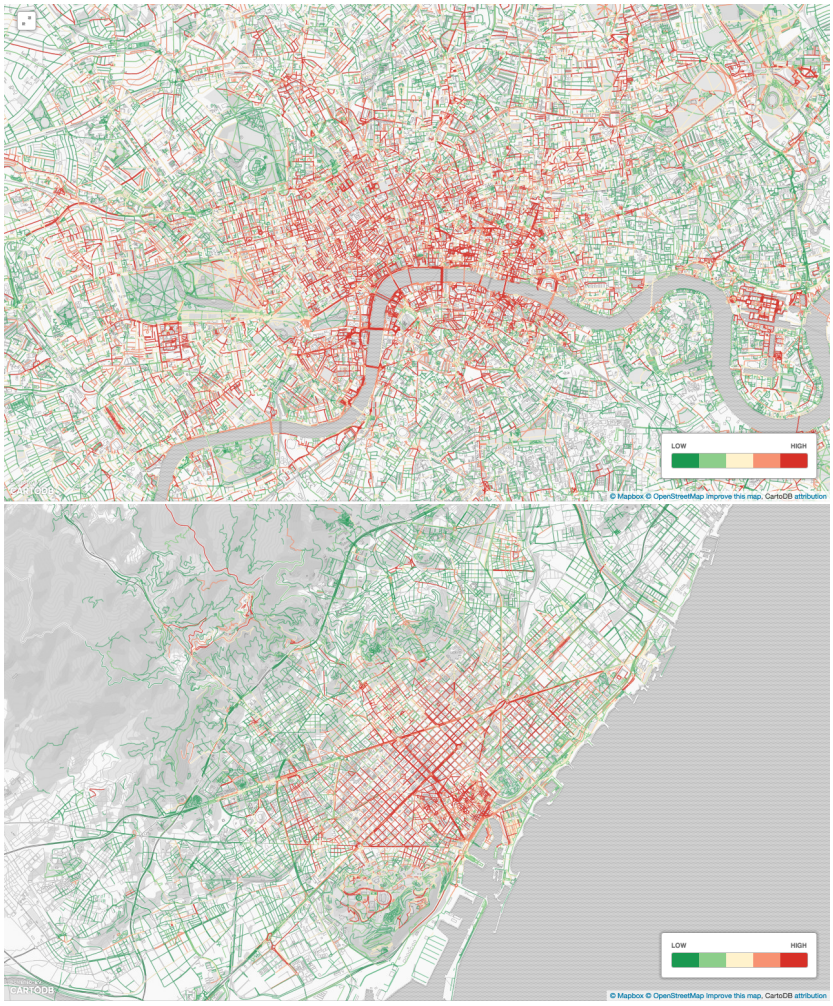

Figure 18: Maps of the diversity of sound tags for each street segment in London (top) and Barcelona (bottom). Only segments with 5 or more tags are displayed.

References

[1] F. Aletta, Ö. Axelsson, and J. Kang. Towards acoustic indicators for soundscape design. In *Proceedings of the Forum Acusticum Conference*, 2014.

[2] F. Aletta and J. Kang. Soundscape approach integrating noise mapping techniques: a case study in Brighton, UK. *Noise Mapping*, 2:50–58, 2015.

[3] F. Aletta, E. Margaritis, K. Filipan, V. P. Romero, Ö. Axelsson, and J. Kang. Characterization of the soundscape in Valley Gardens, Brighton, by a soundwalk prior to an urban design intervention. In *Proceedings of the Euronoise Conference*, page 1–12, 2015.

[4] T. C. Andringa and J. J. L. Lanser. How pleasant sounds promote and annoying sounds impede health: A cognitive approach. *International Journal of Environmental Research and Public Health*, 10(4):1439–1461, 2013.

[5] Ö. Axelsson. How to measure soundscape quality. In *Proceedings of the Euronoise Conference*, pages 1477–1481, 2015.

[6] Ö. Axelsson, M. E. Nilsson, and B. Berglund. A Swedish instrument for measuring soundscape quality. In *Proceedings of the Euronoise Conference*, 2009.

[7] Ö. Axelsson, M. E. Nilsson, and B. Berglund. A principal components model of soundscape perception. *The Journal of the Acoustical Society of America*, 128(5):2836–2846, 2010.

[8] M. Becker, S. Caminiti, F. Donato, L. Francis, P. Gravino, M. Haklay, A. Hotho, V. Loreto, J. Mueller, F. Ricchiuti, V. Servedio, A. Sirbu, and F. Tria. Awareness and Learning in Participatory Noise Sensing. *PLoS ONE*, 8(12), 2013.

[9] R. Beelen, G. Hoek, D. Vienneau, M. Eeftens, K. Dimakopoulou, X. Pedeli, M.-Y. Tsai, N. Künzli, T. Schikowski, A. Marcon, et al. Development of NO_2 and NO_x Land Use Regression Models for Estimating Air Pollution Exposure in 36 Study Areas in Europe–the ESCAPE project. *Atmospheric Environment*, 72:10–23, 2013.

[10] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.

[11] B. Blesser and L. Salter. *Spaces Speak, Are You Listening?: Experiencing Aural Architecture*. MIT Press, 2009.

[12] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10), 2008.

[13] A. Brown, J. Kang, and T. Gjestland. Towards standardization in soundscape preference assessment. *Applied Acoustics*, 72(6):387–392, 2011.

[14] P. Clifford, S. Richardson, and D. Hémon. Assessing the significance of the correlation between two spatial processes. *Biometrics*, pages 123–134, 1989.

[15] W. J. Davies. Editorial to the Special Issue: Applied Soundscapes. *Applied Acoustics*, 2(74):223, 2013.

[16] M. Eeftens, R. Beelen, K. de Hoogh, T. Bellander, G. Cesaroni, M. Cirach, C. Declercq, A. Dedele, E. Dons, A. de Nazelle, et al. Development of Land Use Regression Models for $PM_{2.5}$, $PM_{2.5}$ Absorbance, PM_{10} and PM_{coarse} in 20 European Study Areas; Results of the ESCAPE Project. *Environmental Science & Technology*, 46(20):11195–11205, 2012.

[17] S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.

[18] J. Gulliver, D. Morley, D. Vienneau, F. Fabbri, M. Bell, P. Goodman, S. Beevers, D. Dajnak, F. J. Kelly, and D. Fecht. Development of an Open-source Road Traffic Noise Model for Exposure Assessment. *Environmental Modelling & Software*, 74:183–193, 2015.

[19] J. Halonen, A. Hansell, J. Gulliver, D. Morley, M. Blangiardo, D. Fecht, M. Toledano, S. Beevers, H. Anderson, F. Kelly, and C. Tonne. Road traffic noise is associated with increased cardiovascular morbidity and mortality and all-cause mortality in London. *European Heart Journal*, 36:2653–2661, 2015.

[20] K. Herranz-Pascual, I. Aspuru, and I. Garcia. Proposed conceptual model of environment experience as framework to study the soundscape. In *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, number 8, pages 2904–2912. Institute of Noise Control Engineering, 2010.

[21] B. Hoffmann, S. Moebus, A. Stang, E.-M. Beck, N. Dragano, S. Möhlenkamp, A. Schmermund, M. Memmesheimer, K. Mann, R. Erbel, et al. Residence close to high traffic and prevalence of coronary heart disease. *European Heart Journal*, 27(22):2696–2702, 2006.

[22] H.-P. Hsieh, T.-C. Yen, and C.-T. Li. What Makes New York So Noisy?: Reasoning Noise Pollution by Mining Multimodal Geo-Social Big Data. In *Proceedings of the* 23^{rd} *ACM International Conference on Multimedia (MM)*, pages 181–184, 2015.

[23] International Organization for Standardization. ISO 12913-1:2014 Acoustics — Soundscape — Part 1: Definition and conceptual framework. *Geneve: ISO*, 2014.

[24] J. Y. Jeon, J. Y. Hong, and P. J. Lee. Soundwalk approach to identify urban soundscapes individually. *The Journal of the Acoustical Society of America*, 134(1):803–812, 2013.

[25] P. Kivy. *Sound sentiment: An essay on the musical emotions, including the complete text of the corded shell*. Temple University Press, 1989.

[26] J. Liu, J. Kang, T. Luo, H. Behm, and T. Coppack. Spatiotemporal variability of soundscapes in a multiple functional urban area. *Landscape and Urban Planning*, 115:1–9, 2013.

[27] S. P. Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[28] N. Maisonneuve, M. Stevens, M. E. Niessen, P. Hanappe, and L. Steels. Citizen noise pollution monitoring. In *Proceedings of the* 10^{th} *Annual International Conference on Digital Government Research*, pages 96–103. Digital Government Society of North America, 2009.

[29] C. Meurisch, K. Planz, D. Schäfer, and I. Schweizer. Noisemap: Discussing Scalability in Participatory Sensing. In *Proceedings of ACM* 1^{st} *International Workshop on Sensing and Big Data Mining*, pages 6:1–6:6, 2013.

[30] S. M. Mohammad and P. D. Turney. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3):436–465, 2013.

[31] D. Morley, H. K. De, D. Fecht, F. Fabbri, M. Bell, P. Goodman, P. Elliott, S. Hodgson, A. Hansell, and J. Gulliver. International scale implementation of the CNOSSOS-EU road traffic noise prediction model for epidemiological studies. *Environmental Pollution*, 206:332–341, 2015.

[32] C. Mydlarz, S. Nacach, T. H. Park, and A. Roginska. The design of urban sound monitoring devices. In *Audio Engineering Society Convention 137*. Audio Engineering Society, 2014.

[33] M. E. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.

[34] M. E. Nilsson and B. Berglund. Soundscape quality in suburban green areas and city parks. *Acta Acustica united with Acustica*, 92(6):903–911, 2006.

[35] J. Pennebaker. *The Secret Life of Pronouns: What Our Words Say About Us*. Bloomsbury, 2013.

[36] R. Plutchik. *The Emotions*. University Press of America, 1991.

[37] D. Quercia, L. M. Aiello, R. Schifanella, and A. Davies. The Digital Life of Walkable Streets. In *Proceedings of the* 24^{th} *ACM Conference on World Wide Web (WWW)*, pages 875–884, 2015.

[38] D. Quercia, L. M. Aiello, R. Schifanella, and K. McLean. Smelly Maps: The Digital Life of Urban Smellscapes. In *International AAAI Conference on Web and Social Media (ICWSM)*, 2015.

[39] D. Quercia, N. K. O'Hare, and H. Cramer. Aesthetic capital: what makes london look beautiful, quiet, and happy? In *Proceedings of the* 17^{th} *ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW)*, pages 945–955, 2014.

[40] D. Quercia, R. Schifanella, and L. M. Aiello. The Shortest Path to Happiness: Recommending Beautiful, Quiet, and Happy Routes in the City. In *Proceedings of the* 25^{th} *ACM Conference on Hypertext and Social Media (HT)*, pages 116–125, 2014.

[41] M. Rosvall and C. T. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123, 2008.

[42] J. Salamon and J. Bello. Unsupervised feature learning for urban sound classification. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 171–175, April 2015.

[43] J. Salamon, C. Jacoby, and J. P. Bello. A Dataset and Taxonomy for Urban Sound Research. In *Proceedings of the* 22^{nd} *ACM International Conference on Multimedia (MM)*, pages 1041–1044, 2014.

[44] R. M. Schafer. *The Soundscape: Our Sonic Environment and the Tuning of the World*. Destiny Books, 1993.

[45] B. Schulte-Fortkamp and D. Dubois. Preface to Special Issue: Recent Advances in Soundscape Research. *Acta Acustica United with Acustica*, 92:I–VIII, 2006.

[46] B. Schulte-Fortkamp and J. Kang. Introduction to the Special Issue on Soundscapes. *The Journal of the Acoustical Society of America*, 134(1):765–766, 2013.

[47] I. Schweizer, R. Bärtl, A. Schulz, F. Probst, and M. Mühläuser. NoiseMap - real-time participatory noise maps. In *Proceedings of the* 2^{nd} *International Workshop on Sensing Applications on Mobile Phones*, pages 1–5, 2011.

[48] J. Selander, M. E. Nilsson, G. Bluhm, M. Rosenlund, M. Lindqvist, G. Nise, and G. Pershagena. Long-term Exposure to Road Traffic Noise and Myocardial Infarction. *Epidemiology*, 20(2):272–279, March 2009.

[49] C. Semidor. Listening to a city with the soundwalk method. *Acta Acustica United with Acustica*, 92(6):959–964, 2006.

[50] M. Southworth. The sonic environment of cities. *Environment and Behavior*, 1969.

[51] S. A. Stansfeld, B. Berglund, C. Clark, I. Lopez-Barrio, P. Fischer, E. Öhrström, M. M. Haines, J. Head, S. Hygge, I. Van Kamp, et al. Aircraft and road traffic noise and children's cognition and health: a cross-national study. *The Lancet*, 365(9475):1942–1949, 2005.

[52] The European Parliament. Directive 2002/49/EC: Assessment and Management of Environmental Noise. *Official Journal of the European Communities*, 189(12), June 2002.

[53] E. Van Kempen and W. Babisch. The quantitative relationship between road traffic noise and hypertension: a meta-analysis. *Journal of hypertension*, 30(6):1075–1086, 2012.

[54] G. Watts, A. Miah, and R. Pheasant. Tranquillity and soundscapes in urban green spaces - predicted and actual assessments from a questionnaire survey. *Environment and Planning B*, 40:170–181, 2013.

[55] G. R. Watts, R. J. Pheasant, and K. V. Horoshenkov. Predicting perceived tranquillity in urban parks and open spaces. *Environment and Planning B*, 38:585–594, 2011.

[56] M. Zentner, D. Grandjean, and K. R. Scherer. Emotions evoked by the sound of music: characterization, classification, and measurement. *Emotion*, 8(4):494, 2008.